

March 2015

Science & Technology

REVIEW

Where Computation Innovates with U.S. Industry

Also in this issue:

Advancing Supercomputer Architectures

Assessing Aging Nuclear Weapons

**More Experiments and Stronger Partnerships at
the National Ignition Facility**

About the Cover

The High Performance Computing Innovation Center (HPCIC) is Lawrence Livermore's "anchor tenant" at the Livermore Valley Open Campus (LVOC). Modeled after research parks, LVOC is an innovation hub for unclassified research and development conducted in partnership with private companies, academia, and a wide range of other organizations. HPCIC provides those partners with access to some of the world's most powerful supercomputers to accelerate the technological innovation that underpins America's economic vitality. The artist's image represents the coming together of Lawrence Livermore supercomputing and U.S. industry in bioscience, just one of the many fields where collaboration is taking place.



Cover design: Tom Reason

About S&TR

At Lawrence Livermore National Laboratory, we focus on science and technology research to ensure our nation's security. We also apply that expertise to solve other important national problems in energy, bioscience, and the environment. *Science & Technology Review* is published eight times a year to communicate, to a broad audience, the Laboratory's scientific and technological accomplishments in fulfilling its primary missions. The publication's goal is to help readers understand these accomplishments and appreciate their value to the individual citizen, the nation, and the world.

The Laboratory is operated by Lawrence Livermore National Security, LLC (LLNS), for the Department of Energy's National Nuclear Security Administration. LLNS is a partnership involving Bechtel National, University of California, Babcock & Wilcox, Washington Division of URS Corporation, and Battelle in affiliation with Texas A&M University. More information about LLNS is available online at www.llnslc.com.

Please address any correspondence (including name and address changes) to S&TR, Mail Stop L-664, Lawrence Livermore National Laboratory, P.O. Box 808, Livermore, California 94551, or telephone (925) 423-3432. Our e-mail address is str-mail@llnl.gov. S&TR is available on the Web at str.llnl.gov.

© 2015. Lawrence Livermore National Security, LLC. All rights reserved. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. To request permission to use any material contained in this document, please submit your request in writing to Public Affairs Office, Lawrence Livermore National Laboratory, Mail Stop L-3, P.O. Box 808, Livermore, California 94551, or to our e-mail address str-mail@llnl.gov.

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.



Prepared by LLNL under contract
DE-AC52-07NA27344

Contents

S&TR Staff

SCIENTIFIC EDITOR

Gregory C. Burton

MANAGING EDITOR

Ray Marazzi

PUBLICATION EDITORS

Paul Kotta and Carolin Middleton

WRITERS

Rose Hansen, Arnie Heller,
Caryn Meissner, Ann Parker,
Lanie L. Rivera, and M. H. Rubin

ART DIRECTOR

Tom Reason

PROOFREADER

Arnie Heller

S&TR ONLINE

Denise Kellom and Pamela MacGregor

PRINT COORDINATOR

Charlie M. Arteago, Jr.

S&TR, a Director's Office publication, is produced by the Technical Information Department under the direction of the Office of Planning and Special Studies.

S&TR is available on the Web at str.llnl.gov

Printed in the United States of America

Available from
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Road
Springfield, Virginia 22161

UCRL-TR-52000-15-3
Distribution Category UC-99
March 2015

Feature

3 Maintaining the Nation's Lead in High-Performance Computing

Commentary by Dona Crawford

4 A Hub for Collaborative Innovation

At the Livermore Valley Open Campus, the High Performance Computing Innovation Center is ushering in a new era of knowledge exchange.

Research Highlights

11 Gearing Up for the Next Challenge in High-Performance Computing

Livermore prepares for the next wave of supercomputing systems and the opportunities they create for scientific research.

16 When Weapons Age and Materials React

Livermore's reactive transport model allows researchers to assess material compatibility and chemical transformations in aging nuclear weapons.

20 An Increased Shot Rate at the National Ignition Facility

Upgrades at the world's largest laser set the course to increase the number of experiments and more quickly advance stockpile stewardship.

Departments

2 The Laboratory in the News

23 Patents and Awards

25 Abstract

Volcanic Activity Contributes to Warming Hiatus

An international research team led by Lawrence Livermore atmospheric scientists found that smaller volcanic eruptions have contributed to a “warming hiatus” over the last 16 years. The warmest year on record was 1998. After that, the steep climb in global surface temperatures observed in the 20th century appeared to level off. Scientists had previously suggested that factors such as increased heat uptake by the oceans and weak solar activity could be responsible for the lull in temperature increases.

After publication of a 2011 paper in the journal *Science*, increased volcanic activity was implicated in the warming hiatus. Prior to the 2011 paper, the prevailing view was that only very large eruptions were capable of impacting global climate. Scientists have long recognized that erupting volcanoes cool the atmosphere by expelling sulfur dioxide, which combines with oxygen in the upper atmosphere to form droplets of sulfuric acid. These droplets can persist for many months, reflecting sunlight and lowering temperatures at Earth’s surface and in the lower atmosphere.

The new research, published in *Geophysical Research Letters*, further identifies observational climate signals caused by recent volcanic activity. Positive signal detection supports recent findings indicating that a series of small 21st-century volcanic eruptions deflected substantially more solar radiation than previously estimated. Says Livermore’s Benjamin Santer, lead author of the recent study, “This new work shows that late 20th- and early 21st-century volcanic activity produced discernible signals in observed temperature, moisture, and reflected sunlight at the top of Earth’s atmosphere.”

Contact: Benjamin Santer (925) 423-2253 (santer1@llnl.gov).

X-ray Laser Reveals Protein Structure

Livermore scientists participated in research that captured the highest resolution snapshots of a protein ever taken with an x-ray laser. The experiment, conducted at the Department of Energy’s SLAC National Accelerator Laboratory, revealed how a protein from photosynthetic bacteria changes shape in response to light.

SLAC’s Linac Coherent Light Source (LCLS) generated x-ray laser pulses about a billion times brighter than x rays from traditional synchrotrons. The fleeting pulses allowed researchers to see atomic-scale details of how the bacterial protein, called photoactive yellow protein or PYP, changes within millionths of a second after exposure to light. Crystallized samples of the protein, measuring about 2 millionths of a meter long, were sprayed into the path of the x-ray laser beam. Some of the crystallized proteins were exposed to blue light to trigger shape changes. (The image, courtesy of SLAC, is an artist’s rendering of the protein being exposed to the laser.) The incident x rays produced diffraction patterns as they struck the crystals and were used to

reconstruct the protein’s three-dimensional structure. Researchers then compared the structures of the light-exposed proteins to structures of proteins not illuminated by the blue light. Snapshots taken at different points in time were compiled into detailed movies.

The experiment marked the first time that the LCLS has been used to directly observe a protein’s structural changes at such a high resolution, says Matthias Frank, one of three participating Livermore researchers, along with Mark Hunter and Brent Segelke. Frank says the experimental results demonstrated that x-ray laser crystallography can be used to probe the atomic-scale details of biological molecules important to medicine and pharmacological research. The team’s results were published in the December 5, 2014, issue of *Science*.

Contact: Matthias Frank (925) 423-5068 (frank1@llnl.gov).


Cells Store Metals in “Cupboards”

Lawrence Livermore researchers, together with collaborators from the University of California, Los Angeles, have discovered that some algae cells build an intracellular compartment to store metals and thereby maintain equilibrium. “We don’t understand very well how cells maintain balance when the cell is stressed by metal excess or metal deficiency,” says Livermore researcher Jennifer Pett-Ridge. “By storing the metal in a special intracellular compartment, the cell creates a bit of a pantry cupboard for itself and can better maintain its equilibrium.” How this matchmaking of metals and proteins occurs with precision has puzzled cell biologists.

The researchers shed light into how such pantry cupboards are maintained even under stressful conditions of metal deprivation. The research team studied the copper content in *Chlamydomonas reinhardtii*, a single-cell green alga. The organism’s copper concentration stays relatively constant over orders of magnitude of extracellular copper but hyperaccumulates when *Chlamydomonas* is starved for zinc.

To understand how a single cell builds a metal pantry, the team used high-resolution elemental imaging. The tools included fluorescent metal sensors, transmission electron microscopy, and nanoscale secondary ion mass spectrometry (NanoSIMS), which is housed at Livermore. NanoSIMS was recently equipped with a Hyperion II ion source, which provides a one-of-a-kind capability to image metals with up to 50-nanometer lateral resolution. With these tools, the team discovered that the bulk of copper in zinc-deficient cells is concentrated in the so-called pantry, hidden from cellular sensors. When the zinc starvation stress is relieved, copper is released and is used preferentially over extracellular sources of copper for biosynthesis.

Contact: Jennifer Pett-Ridge (925) 424-2882 (pettridge2@llnl.gov).





Maintaining the Nation's Lead in High-Performance Computing

LAURENCE Livermore National Laboratory (LLNL) has a long and distinguished history in high-performance computing (HPC). Our Laboratory has pushed the state of the art in computing to fulfill missions in national security as well as the basic science underpinning all Laboratory research and development. HPC capabilities enable progress in a diverse set of applications, from elucidating the structure of proteins to optimally operating a smart electric grid to understanding the complex operation of nuclear weapons.

This tradition of innovation continues today. We are preparing to achieve another major advance in HPC to make our current petascale (10^{15} operations per second) physics and engineering simulation models more predictive, to sustain our nation's nuclear weapons stockpile without the need for full-scale underground nuclear tests. At the same time, our Laboratory is expanding partnerships in HPC to invigorate our programs through collaborative research and to enable the U.S. economy to benefit from the full potential of HPC. These expansive activities are the focus of the article beginning on p. 4.

The partnerships being forged at the High Performance Computing Innovation Center (HPCIC), sited at the Livermore Valley Open Campus (LVOC), are crucial not only to our missions but also to our ability to attract the top scientific talent needed to keep our Laboratory on the cutting edge. New exploratory areas include approaches to extracting information from big data coupled with large-scale constitutive modeling and simulation. These efforts greatly benefit U.S. industry and academia, enabling scientific discovery and the development of new products and materials. The results of these interactions feed back into mission simulations, bolster the nation's overall economic competitiveness in the world marketplace, and maintain the nation's global lead in HPC. Therefore, the diffusion of HPC from the Department of Energy (DOE) national laboratories to a broader community is a priority. HPCIC is designed as a hub to stimulate such HPC innovation, providing a research park-like environment for unclassified research and development.

The next major advance in the Laboratory's HPC foundations is a system called Sierra, which will be about five to seven times more powerful than today's most advanced machines. In November 2014, DOE announced the signing of a contract with IBM, in partnership with NVIDIA and Mellanox, to deliver this next-generation supercomputer to Livermore beginning

in 2017. It will be used for the most demanding scientific and national-security simulation and modeling applications and will enable continued U.S. leadership in computing. Acquired by the National Nuclear Security Administration's (NNSA's) Advanced Simulation and Computing program, Sierra will serve as an NNSA-wide resource for stockpile stewardship. Procurement of the system is part of the DOE-sponsored Collaboration of Oak Ridge, Argonne, and Livermore to accelerate the development of advanced HPC.

Sierra is an important step toward exascale (10^{18} operations per second) computing, and its design presents tremendous technical challenges and opportunities. Specific issues include new programming models, the reliability and resilience of the hundreds of thousands of components, and the cost associated with data movement on new memory hierarchies, which need many megawatts of electric power. Resolving this energy issue requires innovations in both hardware and software. The article beginning on p. 11 describes Sierra and other HPC architecture options that are being considered. Livermore is deeply engaged because, to the extent possible, computer architectures must be designed with software requirements in mind—and in this case, the requirements are some of the world's most demanding and are of utmost importance to U.S. security.

HPC has been a central strength of the Laboratory—vital to mission success and a source for scientific discovery and technological innovation. We are taking the next big step to advance HPC and working to make certain that the nation benefits broadly from it.

■ Dona Crawford is associate director for Computation.




A Hub for Collaborative Innovation

LAURENCE Livermore has been a leader in innovative science and technology for more than six decades, working to ensure the nation's security and delivering solutions to other important national needs. Many of the Laboratory's accomplishments result from strategic partnerships with private industries and academic institutions. Breakthrough technologies developed through mission-related research are often transformed into products and capabilities that lead to new industries and economic growth.

To broaden the reach of these partnerships, the Department of Energy's (DOE's)

National Nuclear Security Administration authorized Lawrence Livermore and Sandia national laboratories to develop a collaboration space outside the two laboratories' fenced perimeters. Established in 2011, the Livermore Valley Open Campus (LVOC) is growing into an innovation hub for unclassified research and development activities. The goals for LVOC are to enhance technology transfer from the national laboratories to the private sector, accelerating the pace of industrial innovation, and to expand the laboratories' knowledge base through greater commercial and academic interactions.



At the Livermore Valley Open Campus, the High Performance Computing Innovation Center is ushering in a new era of knowledge exchange.

LVOC is modeled after research parks, which operate with less restrictive security and access controls than are required for the often-classified research at Lawrence Livermore and Sandia. The business and operating principles for the open campus are designed to enhance collaborations with a wide range of organizations, from industrial partners and think tanks to academic institutions and federal, state, and local agencies.

According to Betsy Cantwell, Livermore's former director for economic development (now at Arizona State University), "LVOC bridges the gap

The artist's rendering in the background above shows the facility planned for the High Performance Computing Innovation Center (HPCIC), which will include offices, conference rooms, and collaboration spaces. (Photograph by George Kitrinis)

between businesses and the national labs." She adds, "By working in an unclassified, highly accessible setting, our partners can progress more quickly on new products and bring them to market faster, increasing their profitability." Collaborations also enhance the laboratories' national security missions, allowing researchers to develop

expertise in new areas. “Partnerships strengthen our research capabilities and help us stay at the forefront of science, technology, and engineering,” adds Cantwell. “We gain industries’ knowledge, and they gain economic value.”

The open campus is being built on a 110-acre parcel along the eastern edge of the two neighboring laboratories. (See *S&TR*, March 2011, pp. 22–25.) Each laboratory has established an “anchor tenant”: Sandia’s is the Combustion Research Facility, and Livermore’s is the High Performance Computing Innovation Center (HPCIC). Livermore’s facility provides offices both for LLNL staff and partners, along with conference and classroom facilities. Plans for the next decade include adding conference space, collaboration facilities, and a visitor’s center.

HPC Leads the Way

High-performance computing (HPC) is integral to every research program at Lawrence Livermore, which is home to some of the world’s most powerful supercomputers. HPCIC was thus an appropriate choice as the Laboratory’s first venture at LVOC. “High-performance computing is part of the Laboratory’s DNA,” says Cantwell. The center opened in 2011 with the goal of cultivating HPC-based collaboration and knowledge exchange in a business-friendly environment. Since then, HPCIC has hosted more than 3,000 events and has received more than 25,000 visitors. “It’s become a hub of collaborative activity and the venue of choice for workshops in a broad range of subjects,” says Cantwell.

The Laboratory is renowned for its HPC resources and expertise. “Having these capabilities is vital to our national security missions,” says computational physicist Frederick Streit, the HPCIC director. “We continually push to keep our capabilities at the foremost edge of computational innovation. Simulations help researchers gain answers more quickly and with greater accuracy, allowing them to develop



HPCIC has become a hub of collaborative activity and the venue for various workshops, such as the 24-hour “hackathons” sponsored by Livermore’s Computation Directorate to encourage collaborative programming and creative problem solving by employees and students at LLNL. (Photograph by Meg Epperly.)

a more detailed understanding of the processes and materials they are studying. And that leads to more confidence in the recommendations we pass on to decision makers.”

As a result of that push to expand capabilities, the HPC resources at Livermore and other national laboratories are often more than a decade ahead of many private companies. “It’s not only our supercomputers that offer better performance,” says Streit, “but also the support systems, software, and expertise we develop in trying to fully exploit these computational resources.” (See the box on p. 7.)

HPCIC makes those same benefits available to industrial partners to accelerate their innovation cycles. Collaborative teams include experts from all organizations involved in a collaboration, whether a private company, academic institution, or the Laboratory. Research topics range from simulations for optimizing a “smart,” or interconnected, electrical grid system and for predicting the availability of renewable energy sources to examining laser–plasma

interactions and discovering new drugs and treatment options to improve human health.

“There’s a growing consensus worldwide that HPC is a key to accelerating the technological innovation that underpins a nation’s economic vitality,” says Streit. For example, complex models now allow developers to create virtual prototypes of new devices and work through design iterations. Simulations might reveal a design flaw or show how a workflow can be improved, allowing developers to implement new ideas before building with materials.

Removing the Hurdles

Streit notes, however, that private industries may hesitate to adopt HPC systems because of the investment required. “Supercomputers are not ‘plug-and-play’ devices,” he says. “To fully exploit a machine’s capabilities, companies need the expertise to design codes that run efficiently on massively parallel systems.” Computational scientists at Livermore have extensive experience working with HPC systems and have developed a rigorous

process for validating the accuracy of the simulated results. By collaborating through HPCIC, private-sector businesses have access to the Laboratory's knowledge, experience, and unclassified supercomputers to pursue new technologies and manufacturing capabilities.

The selection process at HPCIC ensures that projects match Livermore strengths and industry needs. "Projects must also help us advance the Laboratory's capabilities," says computer scientist Deborah May, who manages the center's business development efforts. May connects promising proposals with the Livermore researchers best suited for those projects. Partnerships are then structured under a formal arrangement, such as a cooperative research and development agreement or DOE's Strategic Partnership Projects, formerly known as Work for Others. Team members then work together for the project's duration, focusing their expertise toward finding innovative solutions to complex challenges.

Building a Knowledge Pipeline

HPCIC also fosters long-term strategic partnerships in research areas that will provide value to Laboratory programs and corporate entities. One such partnership, called Deep Computing Solutions, expands on Livermore's 20-plus-year relationship with IBM in developing HPC systems for stockpile stewardship. Through this effort, computational experts from IBM and Livermore work with collaborators from U.S. industries to accelerate the development of new technologies that will benefit both the nation's security and its economy.

A partnership with RAND Corporation focuses on developing new capabilities in scalable policy analytic methods. "The RAND partnership with Livermore offers the opportunity to pursue new understanding and potential solutions to even the most intractable current and future policy problems," says Susan Marquis, RAND vice president for Emerging Policy Research and Methods and dean of the Pardee RAND Graduate School.

A Complete Ecosystem of High-Performance Computing Resources

The High Performance Computing Innovation Center (HPCIC) provides industrial partners access to Lawrence Livermore's unmatched computing resources. Frederick Streitz, the HPCIC director, says, "A massive supercomputer with one and a half million cores is impressive, but by itself, it's of limited use. It's the Laboratory's computational ecosystem—our hardware, software, and experienced people—that enables many discoveries."

At Livermore, the high-performance computing (HPC) "ecosystem" includes more than 25 systems for parallel numerical simulations, visualization, and data analytics backed by several massive parallel file systems and storage archives. The Laboratory also has comprehensive software assets for HPC, including numerical libraries, highly scalable scientific application codes, and the supporting system software and tools. "But most importantly," says Streitz, "we have the necessary team of talent and experience, from system administrators to experts in performance optimization, applied mathematicians, and computer scientists."

The Catalyst supercomputer is available for exploring data-intensive technologies, architectures, and applications. Developed by a partnership of Cray, Intel, and Lawrence Livermore, this Cray CS300 machine is a resource for the National Nuclear Security Administration's Advanced Simulation and Computing Program. The Catalyst cluster has 7,776 processors (or compute cores) distributed over 324 nodes, each with nearly a terabyte (1012 bytes) of addressable memory. The nonvolatile memory (or NVRAM) retains files even when the power is off, as on a USB memory stick or an MP3 player. Catalyst allows researchers to explore new approaches to big data analytics and hierarchical memory systems.

For research that requires advanced architectures, industrial collaborators can access Vulcan, an IBM BlueGene/Q system able to process 5 quadrillion flops (or petaflops). It came online in 2013 and consists of 24 racks with a total of 24,576 compute nodes, or 393,216 compute cores. Vulcan's architecture is identical to that of Sequoia, the 20-petaflop machine dedicated to classified national security research.

"Our broad, robust ecosystem keeps us at the vanguard of computing," says Streitz. "With it, we have both a long history and a promising future of turning compute cycles into science and solutions."



HPCIC Director Frederick Streitz (left) meets with Doug East, the Laboratory's chief information officer, in front of the Laboratory's Vulcan supercomputer. This IBM BlueGene/Q system can process 5 quadrillion floating-point operations per second. (Photograph by Laura Schulz.)

Says Jim Brase, deputy associate director for big data in Livermore's Computation Directorate, "Although there are policy elements to the science work we conduct, the Laboratory doesn't have the extensive experience in policy that RAND does. Coming together helps us gain knowledge and establish new expertise in an area that is relevant to our mission work."

In the initial project with RAND, researchers revisited an earlier study that evaluated water management strategies for the Colorado River Basin. The joint team used the University of Colorado Boulder's RiverWare and RAND's Robust Decision Making analytical framework on a Laboratory system to model the effects of the policy options considered in the original study as well as additional strategies, producing results within hours. "The RAND simulations demonstrate well how large-scale computation can change the nature of the game," says Streitz. "Developing the ability to explore complex data sets and decision options at a scale previously impossible could revolutionize the way decision makers, policy analysts, and the research community approach some of today's most challenging issues."

Another area of interest for HPCIC involves independent software vendors, which develop and sell the scientific and engineering codes that dominate American industry. "Companies have built entire workflows around these commercial codes," says Streitz, who also chairs a software working group for the Council on Competitiveness. "Unfortunately, their codes don't scale to the newer generations of computational platforms available to industry." HPCIC collaborations with vendors are focused on software evolution to scale, expand, and optimize these codes so they will run on current and future HPC systems. HPCIC provides access to the hardware needed for testing and validating the codes. Collaborators work closely with Laboratory scientists who have the experience required to program at scale, allowing them to more quickly modify their codes.

"Just as our computers are a decade ahead of most of industry, the same can be said of our codes and our programming knowledge," says Streitz. "We want to help software vendors transition to more powerful computing platforms. If they move forward, industry will follow."

Collaborations Across Industries

HPCIC partnerships also offer access to Livermore research centers such as the Turbulence Analysis and Simulation Center (TASC). Managed by the Engineering Directorate, TASC applies advanced numerical methods to simulate and analyze turbulent mixing and reactive flows. These simulation capabilities are important in research to design jet, rocket, and internal combustion engines. Reactive flow models can also be used to predict weather patterns, the atmospheric dispersion of pollutants, and the availability of wind power for electricity generation.

In 2012, engineers from TASC collaborated with General Electric Global Research on simulations to improve the efficiency of jet engines as part of the Laboratory's HPC4energy initiative. Under HPC4energy, six U.S. energy industries partnered with Livermore computational scientists to demonstrate the potential of advanced computing to provide solutions to energy and environmental challenges. In addition to General Electric Global Research, HPC4energy projects involved GE Energy, Potter Drilling, United Technologies, ISO New England, and Bosch. (See *S&TR*, June 2012, pp. 24–25.)

In an HPCIC strategic partnership, Livermore researchers are working with colleagues from RAND Corporation to explore the use of HPC applications for analyzing options in public policy. At this conference, RAND and Lawrence Livermore held a demonstration for water managers—including the U.S. Department of Interior's Bureau of Reclamation and the California Department of Water Resources—and others to show how Laboratory HPC could perform a study of the Colorado River Basin in hours instead of the several days required initially, while evaluating additional strategy options at the same time. (Photograph by George Kitrinos.)



Eugene Litinov, senior director of Business Architecture and Technology at ISO New England, noted that access to the Laboratory's computing capabilities allowed his company to "think differently about problems. You can ask questions you didn't think of asking before."

In another energy-related collaboration, Livermore and IBM computational experts helped Energy Exemplar Corporation to simulate the nation's electric power grid. This research led to a massively parallel implementation of Energy Exemplar's energy market simulation software, increasing the software's performance a thousandfold. "Our work with Energy Exemplar served as a demonstration of the capability that can be developed by applying HPC to the challenges of our energy system," says Streitz. "Ensuring the efficiency, reliability, and safety of the energy grid are national security challenges."

A partnership with Cymer Corporation is applying codes originally developed for experiments at the National Ignition Facility to help the company develop an extreme ultraviolet (EUV) light source. The Cymer light source will use laser-heated tin droplets to generate EUV light for creating computer circuits on silicon wafers. The collaborators are working to shorten the time required to develop the new light source, which will enable chip manufacturers to increase the number of transistors on a computer chip and thus significantly improve the performance of future supercomputers.


"Using our codes to simulate the conditions in Cymer's experiments resulted in a more robust and broadly applicable modeling capability for us," says Livermore computational physicist Steve Langer, who led the Cymer project. "We also gain a lot of satisfaction from applying our expertise to activities that have a direct impact on everyday life."

Oil and natural gas service provider Baker Hughes and Livermore geoscientists are teaming to better predict the formation of hydraulic fractures for shale oil and gas

production. A new supercomputer code being developed will include the important physical processes that control fractures in rocks. In a related project, the Laboratory is partnering with a company to speed up simulations of fluid flow in oil and gas reservoirs, which will improve business decisions for oil and gas development.

Biomedical Payoffs

Biomedical research also benefits from the incorporation of HPC into the discovery process. Under the Laboratory's Medical Countermeasures Program, Livermore scientists and colleagues from a pharmaceutical company ran HPC codes to accelerate development work on a new antibiotic to treat infections. Simulations helped researchers screen new compounds, create detailed maps of target proteins that drugs will bind to, and predict the physiochemical properties of the compounds as well as potential side



The Cardiod code, developed by a team of IBM and Lawrence Livermore computational scientists, replicates the electrophysiology of the human heart. This snapshot from a Cardiod calculation shows a heart immediately following a heartbeat, during which electrical excitation travels through the heart's cells. The recovery of cells to their resting voltage (blue) from this excited state (red) varies from region to region. (Image courtesy of IBM.)



Livermore's Catalyst supercomputer, designed for data-intensive computing, is available for research collaborations with U.S. industries and academic institutions. (Photograph by Laura Schulz.)

effects. This collaboration resulted in a patent application for a broad-spectrum antibiotic that was developed in only three months. Modeling has also proven effective for designing new drug delivery platforms and developing vaccines and antimicrobials.

A biomedical effort between IBM and LLNL produced *Cardioid*, the most realistic simulation of a beating human heart ever developed. (See *S&TR*, September 2012, pp. 22–25.) The supercomputer code driving *Cardioid* models the heart's electrical system, the current that originates from and travels through the heart and causes it to beat and pump blood. This breakthrough simulation capability holds promise in helping researchers better understand how the heart works and responds to different medicines—important information for discovering new drugs and patient-specific therapies to treat cardiovascular disease and improve heart health.

Data-Intensive Computing

An important new area for partnerships focuses on data-intensive computing, in which data ranging from terabytes to petabytes (10^{12} – 10^{15} bytes) in volume are processed. “We believe that advancing ‘big data’ technology is a key to accelerating innovation,” says Brase. “Over the next decade, global data volume is likely to grow to more than 35 zettabytes—that’s 35 trillion gigabytes. We want to figure out how to extract value from this wealth of raw information so we can better inform decision makers.”

HPCIC is expanding the number of partnerships involving data-intensive science and analytics, with simulations performed on Livermore's Catalyst supercomputer. With Catalyst, researchers can store enormous reference databases in

memory and run expansive analyses with more detail and resolution than previously available. Example projects include analyzing the genomes of disease-causing microbes, extracting predictive health indicators from a hospital database, and testing complex models to improve video searches.

Looking Ahead

According to Cantwell, LVOC represents a “sea change” for Lawrence Livermore. “HPCIC has expanded the breadth of collaboration opportunities for Lawrence Livermore while providing tangible benefits to companies.”

She adds that the Laboratory is benefitting from industry as well, often in areas that are central to national security, such as cybersecurity, energy generation, advanced manufacturing, and bioscience. “Collaborations with Bay Area industries help us stay abreast of developments in these rapidly advancing fields.”

Camille Bibeau, assistant to the Laboratory's director of economic development, is exploring ways to revitalize and expand Lawrence Livermore's open-campus infrastructure. For example, she is working with the University of California (UC) to renovate Hertz Hall—formerly home to the Department of Applied Science at UC Davis—to include UC-wide programs in areas of mutual interest.

Bibeau is also exploring potential third-party financial arrangements to build a permanent home for HPCIC, which is currently housed in a temporary modular building. DOE approval is pending on a proposed 100,000-square-foot facility targeted for completion in 2019. The new facility will be eight times larger than the temporary facility and accommodate up to 400 tenants. As HPC collaborations continue

to grow, scientists and engineers from industry and academia will benefit from working side-by-side with Livermore staff.

“Challenges include presenting low-risk business models that DOE will accept for building new infrastructure on campus,” Bibeau says. Under one proposed arrangement, DOE would lease land to a third party, which would build the facility. The Laboratory would then lease the facility and over time provide space to its strategic partners. With new facilities in operation, Bibeau foresees LVOC hosting international conferences and workshops and expanding educational programs for local communities and for the nation.

Some of the strongest support for LVOC comes from businesses and governments in the Tri-Valley area, which includes the cities of Livermore, Pleasanton, Dublin, San Ramon, and Danville. In that vein, Bibeau is working closely with the Tri-Valley's startup incubator i-GATE, one of twelve Innovation Hubs (or i-Hubs) designated by the State of California to foster the growth of technology-oriented companies. Bibeau helped found i-GATE, which is funded by Tri-Valley governments and private donations.

“Interest in collaboration opportunities with Lawrence Livermore is increasing,” says Bibeau. With LVOC and HPCIC, the Laboratory is developing an innovation pipeline that will contribute to national security and U.S. economic competitiveness.

—Arnie Heller

Key Words: data-intensive computing, *Cardioid* code, Catalyst supercomputer, high-performance computing (HPC), High Performance Computing Innovation Center (HPCIC), Livermore Valley Open Campus (LVOC), Vulcan supercomputer.

For further information contact Frederick Streitz (925) 423-3236 (streitz1@llnl.gov).

Gearing Up for the Next Challenge in High-Performance Computing

EVERYTHING changes. Nowhere is that maxim more apparent than in the world of computing. From smartphones and tablets to mainframes and supercomputers, the system architecture—how a machine's nodes and network are designed—evolves rapidly as new versions replace old. As home computer users know, systems can change dramatically between generations, especially in a field where five years is a long time. Computational scientists at Lawrence Livermore and other Department of Energy (DOE) national laboratories must continually prepare for the next increase in computational power so that the transition to a new machine does not arrest efforts to meet important national missions.

That next jump in power will be a big one, as new machines begin to approach exascale computing. Exascale systems will process 10^{18} floating-point operations per second (flops), making them 1,000 times faster than the petascale systems that arrived in the late 2000s. Computational scientists will need to address a number of high-performance computing (HPC) challenges to ensure that these systems can meet the rising performance demands and operate within strict power constraints.

Up the Supercomputer Highway

This is not the first sea change presented by advances in supercomputing. Since the first computers arrived in the 1950s,

four eras have made an entrance, each with its advantages and challenges. In the mainframe era, large sequential processing machines executed computer code instructions one at a time, in serial fashion. Memory capacity (the amount of data that could be stored) was often an issue for mainframe computers, limiting the size of applications and requiring developers to find a balance between memory usage and application.

The vector era of the 1970s and 1980s offered a large performance boost. With vector processors, computers could gather sets of data elements scattered around the system's memory and align them into vector registers, where codes could efficiently operate on the data and send the results back into memory. This architecture mapped favorably to scientific programs, where arrays of data with different values to be computed by the same set of instructions could now be processed concurrently. Ultimately, researchers found they could vectorize only about 30 percent of the operations performed by Livermore's most complex national security multiphysics codes. Therefore, to improve overall runtimes, Laboratory scientists and computer architects at the partnering vendors worked together to improve the scalar performance of serial (one-at-a-time) operations that could not be vectorized. They also continued to work on vectorizing codes to improve performance even further.

Vector processing gave way to the distributed-memory era in the 1990s, when commodity serial processors connected by fast networks proved to be a cost-effective architecture. Algorithms were again redesigned for parallel programming, using message-passing routines for efficient communication between nodes. The boost in performance came from parallelization across nodes and from increases in scalar performance on the processors.

To further improve performance and overcome a growing gap between compute and memory speeds, developers added a small amount of fast memory (called a cache) inside each processor. Cache keeps data close to the central processing unit (CPU) and available for reuse, eliminating extra operations to store and fetch data from main memory. Unfortunately, the memory capacity per core and the memory bandwidth between cores and local memory have not kept pace with increases in peak floating-point performance, creating ever more serious choke points for applications.


Attempts to address this issue require innovation in both hardware and software, leading to the fourth HPC era: many-core computing. This architecture is typified by either a very large number of CPU cores on a node, or an accelerator—often graphics processing units (GPUs) such as those originally developed for three-dimensional (3D) rendering in video games. The node design can also include complex memory hierarchies. For example, one section of main memory can be fast and small, and the other is large but slow. Livermore’s Sequoia supercomputer is a harbinger of such advanced architectures, built with a large number of

low-powered cores, yet retaining a “flat” memory hierarchy within a node. An identifying trait of the many-core era is a requirement to shift to threaded processes, again requiring radical algorithm redesigns for the codes and continued innovations in languages and compilers.

Data Movement and Parallelism

Livermore computational physicist Bert Still explains how the next-generation HPC systems will affect the current situation. “In the past, applications were developed on systems where the main work of computing—floating-point operations—took place on the CPU,” says Still, the deputy project leader for the Advanced Architecture Software Development project funded by the National Nuclear Security Administration’s Advanced Simulation and Computing (ASC) Program. “We and our industrial partners focused on streamlining this work in both applications and computing architectures.” As a result, data packets and streams were often directed around the computer system—in and out of memory and various subsystems—with little regard for the electricity required to move that memory around the machine. Now that more data must be stored, handled, and manipulated, the electrical cost of moving data could prove prohibitive. Thus, the first challenge is to reduce data motion, either by designing algorithms and applications that perform as many calculations as possible on a piece of data before returning it to main memory, or by minimizing the communication required with neighboring nodes.

Each era of high-performance computing brings its own challenges along with increased performance capabilities. The first three—the mainframe, vector, and distributed-memory eras—have passed. In the current many-core era, node designs deploy many central processing units (CPU) cores or a graphics processing unit (GPU) accelerator with various memory configurations. Livermore’s Sequoia stands on the threshold of this era.





Livermore's Sequoia represents another leap in the evolution of supercomputer architectures—a many-core configuration combined with a node design featuring a “flat” memory hierarchy. Even with its large number of cores, Sequoia is significantly more energy efficient than a conventional computer system thanks to the low power required by these cores. Ranked as the world's fastest supercomputer for a time, this machine can run suites of uncertainty calculations used to increase confidence in the predictions of computer models. (Photograph by Bob Hirschfeld.)

Still notes that although Sequoia is significantly more energy efficient than a conventional computer system, it consumes 9.6 megawatts at peak speed. “If 1 megawatt costs \$1 million per year, you can see how the costs push us toward energy-efficient advanced architectures,” says Still. “If the architecture and codes stayed the same and we just pushed to a bigger system, the power requirements would be prohibitive. The annual electric bill for running that system could be several hundred million dollars—far more than the cost of the capital equipment.”

The second challenge involves the increased parallelism in the system as computer architects design machines for yet more performance. In the past, performance gains were accomplished by pushing the clock speed (the rate at which each microprocessor executes instructions) and adding power-hungry complexity (more transistors) to CPUs to automatically exploit low-level parallelism. “The ‘good old days’ of increasing clock rates ended nearly a decade ago,” says Livermore scientist Rob Neely, who leads the Advanced Architecture Software Development project. “We now redeploy those extra transistors in multicore CPUs to boost overall performance.”

Possible Answers



According to Still, a radical shift in architectures is required to minimize data motion and further reduce computational time. One approach is to design cores and memory within each node in a way that increases parallelism and concurrency. “We already see this trend in successive generations of the IBM BlueGene architectures over the last decade,” says Still. “In 2005, the BlueGene/L machine had 196,608 cores in 98,304 nodes. By 2012, Sequoia had 1.6 million cores in the same number of nodes.”

The BlueGene architecture relied on a homogeneous node consisting of multiple, identical cores. A competing architecture uses a heterogeneous node that combines GPUs with commercially available high-performance CPUs. GPUs have hundreds of cores that handle thousands of software threads simultaneously. They can take gigabytes of data and repeat the same operations very quickly by using thousands of streaming processors. Calculations that cannot effectively use GPUs are processed by CPUs instead.


Future heterogeneous designs will lessen the burden on the programmer by allowing the distinct memory between CPU and GPU to appear as a single unified memory. Explicitly managing data movement between CPU and GPU will no longer be required. However, to gain the best performance, developers will need to optimize the application by providing ample “hints” to the compilers indicating where data should be placed.

Another advanced architecture is one that is similar to the BlueGene supercomputers but works with both fast and slow memory in a configuration called nonuniform memory access (NUMA). The small, fast memory with high bandwidth is located on a many-core package. The large, slow memory is farther away and accessed by a slower link. In the NUMA configuration, each core has an instruction stream and fetches its own data but may share a cache with others cores on the chips.


The processing-in-memory architecture, which adds a simple arithmetic unit in or near main memory, is yet another design being considered. This approach would eliminate some traditional data motion, such as transferring data arrays to CPUs for calculations and returning the results back to memory for storage. Instead, a CPU could simply issue an instruction to the memory subsystem



1995
Single CPU per node with
main memory



2000–2010
Multiple CPUs per node
sharing main memory



2000–2010
Accelerators usher in era
of heterogeneity

to return the sum of that array. “In this design,” says Neely, “a subset of the operations is offloaded to the memory processor, further reducing data motion and memory bandwidth requirements between the main CPU and memory.”

All of these architectures include new memory technologies, and the field is evolving rapidly still. ASC leaders are evaluating candidate architectures with the goal of acquiring the best performance gain possible with the fewest modifications to the million-plus lines of code in the multiphysics packages. “We need computer programs that can express the actions we want and a system to perform in languages such as C++, Python, and FORTRAN,” says Neely. “To get the necessary performance gains, we must focus on the whole picture: hardware, software, and applications.”

As Still points out, complex science questions are looming, and they involve calculations that current machines cannot handle. Whether it’s simulating the interactions of intense laser beams with plasmas, the atomic-level behavior of metals under extreme stress and strain, or the effects of local weather variability on global climate systems, the more accurately simulations can mimic and predict natural processes, the better. Improving predictive capability involves more data, more processing power, and more complex calculations. Given the current flux in computer architecture design, scientists face the challenge of rethinking or even rewriting codes to ensure confidence in the modeled predictions.

Working Together for Success

Even as the experts peer into the future, the Collaboration of Oak Ridge, Argonne, and Lawrence Livermore national laboratories (CORAL) is focusing on the next big near-term system. In January 2014, CORAL announced a joint request for

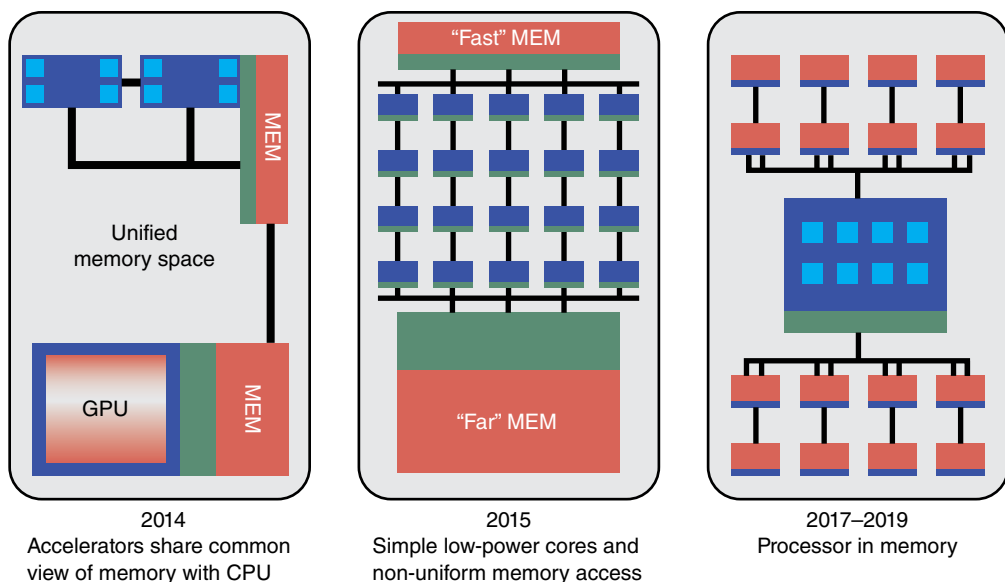
proposals for next-generation supercomputers that offer peak performance of at least 100 quadrillion flops (petaflops), about 5 times the capability of Sequoia but only 10 percent of the exascale mark. Under CORAL, scientists at the three laboratories are working with vendors to develop computer systems that will be deployed in 2017 and 2018. Livermore’s system will be used for national security calculations to support nuclear stockpile stewardship under the ASC Program. Oak Ridge and Argonne will use their supercomputers to perform missions for DOE’s Office of Science, under the Advanced Scientific Computing Research Program.

Bronis de Supinski, chief technology officer for Livermore Computing, explains, “Our collaborative goal was to choose two systems that, as a set, offer the best overall value to DOE. We want diversity of technologies and vendors as well as systems that will provide value to the Office of Science laboratories.”

On November 14, 2014, Secretary of Energy Ernest Moniz announced that IBM, working closely with OpenPOWER Foundation partners NVIDIA and Mellanox, was chosen to design and develop systems for Lawrence Livermore and Oak Ridge. The design uses IBM Power architecture processors connected by NVLink to NVIDIA Volta GPUs. NVLink is an interconnect bus that provides higher performance than the traditional peripheral component interconnect for attaching hardware devices in a computer, allowing coherent direct access to GPU and memory. The machine will be connected with a Mellanox InfiniBand network using a fat-tree topology—a versatile network design that can be tailored to work efficiently with available bandwidth.

IBM will initiate delivery of the Livermore machine, called Sierra, in 2017. Sierra will provide more than 100 petaflops of capability. “We estimate that the peak power required to run this machine will be about 10 megawatts—just slightly higher than

New programming models required



HPC nodes have evolved over the decades, becoming ever more complex and densely packed. In 1995, machine architecture featured a simple node with a single CPU and a small cache for storing copies of frequently used data from the main memory. Current machines have multicore CPU-cache units that share a common main memory.

Sequoia,” says de Supinski. A small, early-access system scheduled for delivery in 2016 will have an earlier generation of the IBM Power processor architecture, NVIDIA Pascal GPUs, and a version of NVLink. “It will be a complete precursor system,” de Supinski adds, “so we can explore the capabilities and begin to deploy some early software and applications on the machine.”

Before Sierra arrives, scientists in the Computation Directorate will work with the vendors to ensure that “no code is left behind when Sierra goes live,” says Michel McCoy, the ASC program director at Livermore. “Having the hardware on the floor is only part of the challenge. We also need system software that boosts the machine’s usability so that applications and key libraries will run efficiently and effectively—not only on Sierra’s massively parallel, accelerator-based nodes but also on alternative architectures and future systems, as well.”

As part of this collaboration, code developers will analyze and modify algorithms, investigate new data structures and layouts, and map workflows onto the new system. Once the early-access system is live, vendors will provide customized training to the Laboratory’s applications scientists, working onsite to share their expertise. “This kind of collaboration allows us to tune our mission-critical application codes and quickly resolve issues as they arise,” says Neely.

McCoy notes that efforts to get the weapons codes ready for Sierra will also benefit the codes that run on Livermore’s unclassified systems. “It’s not just stockpile stewardship that depends on HPC capabilities,” he says. “We have a wide array of projects that rely on our supercomputing resources, from biomedical research to climate modeling and energy production.”

The Laboratory’s Multiprogrammatic and Institutional Computing (M&IC) Program, led by Brian Carnes, brings tailored, cost-effective unclassified computing services to all

Livermore programs and scientists. “Through M&IC, we buy a smaller version, or ‘clone,’ of the larger system purchased for the ASC Program,” says Carnes. “This strategy ensures that all of the Laboratory’s science and technology areas have up-to-date computational resources. It’s also more efficient if researchers across the Laboratory can work in a homogeneous computing environment, whether their projects are classified or unclassified.”

Stepping into the Future

Still and others are looking forward to the increased capability that Sierra will bring. “On Sequoia, we can run suites of large 2D or small 3D uncertainty calculations, which are used to validate the computer models,” says Still. “Sierra will allow us to do moderate to large 3D uncertainty calculations. It’s another step up in our capabilities to run these complex problems.”

DOE’s support for HPC brings together the people who build the machines, those who write the codes, and those who use the software and hardware to explore important questions in science. The speed with which computing technology changes presents exciting opportunities while introducing challenges. “The problems may seem daunting, but they can be solved,” says Still. “We know exascale won’t be the end, and we want the Laboratory to be ready to address those issues when they arise.”

—Ann Parker

Key Words: Advanced Simulation and Computing (ASC) Program, BlueGene, Collaboration of Oak Ridge, Argonne, and Lawrence Livermore national laboratories (CORAL), central processing unit (CPU), graphics processing unit (GPU), high-performance computing (HPC), nonuniform memory access (NUMA) configuration, processing-in-memory supercomputer architecture, Sequoia, Sierra.

For further information contact Bert Still (925) 423-7875 (still1@llnl.gov).



When Weapons Age and Materials React

PRICELESS works of art, such as the Mona Lisa, and nuclear weapons systems rarely have anything in common. But as it turns out, research into how materials in close contact in a sealed environment react with each other and change can be important for both art preservation and stockpile stewardship.

Ensuring a safe, reliable, and secure nuclear deterrent requires scientists to understand weapons performance and the technical issues related to how these systems age. To more closely examine the fundamental chemical transformations that contribute to component aging, Livermore scientists Elizabeth Glascoe, Yunwei Sun, and Stephen Harley have developed a reactive transport model for assessing the compatibility and chemical kinetics of materials inside nuclear weapons systems.


Reactive transport codes are typically used to model chemical reactions and fluid transport through a medium. “Reactive transport is important in any application where materials in contact with one another react and change over time,” says Glascoe. For stockpile stewardship applications, the codes need to model gas sorption as opposed to fluid transport. “In aging weapons assemblies, the chemical compatibility of the materials is an integrated process of outgassing, sorption, mass transport, and chemical reaction kinetics,” says Glascoe.

With funding from Livermore’s Laboratory Directed Research and Development Program, Glascoe and her team have produced a novel computational code and experimental capability that can predict material aging, compatibility, and headspace gas composition for any geometric assembly of materials. “For our purposes, simple diffusion models would have been too rudimentary,” says Glascoe. “We had to develop more advanced mathematical models that incorporate sorption, diffusion, and chemical kinetics to achieve these results.” The team’s model will better predict the long-term behavior of weapons materials and allow scientists to develop more robust system components and nondestructive surveillance capabilities for managing the stockpile.

A Three-Pronged Approach

In any sealed container, including weapons systems, the headspace is the area at the top of the vessel that is not occupied by a solid. If multiple materials housed in the same container begin to react with one another, the headspace can fill with gases produced as the chemicals comprising those materials transform. The new gases may then attack the materials, causing them to degrade more quickly. Over time, what began as seemingly minor material incompatibility could result in substantial damage to the entire system. “Our goal was to develop a predictive model that answered three main questions: where are the gases in the headspace from, why are they there, and are they a problem?” says Glascoe. “By analyzing the headspace gas spectrum, we can identify what are considered normal conditions for a weapons system.”

Using the MATLAB code, the researchers configured their own mathematical and physics models to simulate reactive transport and chemical reaction kinetics while accounting for sorption—how substances become attached to another—and transport—how the materials diffuse through one another. The reactive transport model simulates different sorption mechanisms and incorporates physics-based finite-element models, uncertainty quantification,



Livermore’s reactive transport model predicts material incompatibilities and material outgassing in sealed assemblies, such as a nuclear weapon. In validation experiments, materials are layered in a sealed vessel and aged. Moisture (H_2O) in the base layer diffuses through the assembly, reacting with other materials to produce gaseous products such as carbon dioxide (CO_2) and acetic acid (HOAc). (Rendering by Kwei-Yu Chu.)

and sensitivity analysis techniques. One equation is used to model absorption, via Henry's Law, in which a gas is taken in by another material and fills in that material's void space. Another equation accounts for Langmuir adsorption, wherein the gas adheres to a material's surface.

"When the concentration of gas molecules is high, Langmuir or Henry sites may nucleate the growth of larger multilayer clusters that cause the gas molecules to pool," says Glascoe. "We developed a third equation to model pooling concentrations." Together, the three equations help the team visualize the contributions of each sorption mechanism and the kinetic exchange between them.



When running the models, the team used a Laboratory-developed software toolkit called PSUADE (Problem-Solving Environment for Uncertainty Analysis and Design Exploration) to perform uncertainty and sensitivity analysis. PSUADE works well for models with a large number of parameters and complex constraints, such as in this study, which required more than

10,000 simulations to determine the optimal parameters and parameter sensitivities for one material. By integrating PSUADE into the model, the team determined the sensitivity of eight parameters, including diffusivity, desorption rate, and porosity for the materials tested.


Working in Tandem

The researchers conducted experiments concurrently with the model development efforts. "We decided to build the capabilities together so we could ensure a good bridge between experiments and modeling and verify the code's accuracy as we progressed through the study," says Glascoe.

Initial experiments characterized such variables as sorption and transport parameters as a function of relative humidity and temperature. The team started with a quasi-one-dimensional (1D) sample material in a temperature-controlled chamber. "We systematically adjusted the relative humidity and monitored the sample's mass changes over time," says Glascoe. "If the sample



Three sorption mechanisms (Henry's Law absorption, Langmuir adsorption, and pooling) dictate the rate of chemical uptake and outgassing in, for example, a silica-filled poly(dimethylsiloxane) material. Understanding these mechanisms is critical to developing a reactive transport model that can predict component aging and material compatibility accurately in a system. (Rendering by Kwei-Yu Chu.)



Experimental results are used to parameterize the mathematical model and characterize the mechanisms of sorption.

gained mass, then moisture was being sorbed into the material. If the sample lost mass, the process was desorptive.”

Experiments were performed over a wide range of humidities and temperatures. The time-dependent behavior of the material provided information about the sorption and transport properties. With the input, researchers were able to use the physics-based mathematical model to determine the sorption and transport parameters and mechanisms. They then used the parameterized model to predict material response under different conditions and different geometries. Data from the 1D experiments will be extrapolated to build a 3D model.

Results from the validation experiments are being compared to model predictions to ensure the model’s fidelity. For the experiments, the team layered different materials inside tightly sealed vessels. Water was the base vapor because it is an aggressive reactant. Materials were chosen based on their reactivity to water vapor and the time for those reactions to occur. Glascoe says, “We wanted materials that would react with the water vapor in a fairly short time and produce outgassing signatures.”

During the experiments, tiny sensors measured the relative humidity and temperature inside the containers. The parameterized model also simulated the validation experiments. When a model’s predictions match experimental results, researchers have confidence in using the model to examine more complicated assemblies, such as a weapon. “Our preliminary validation experiments showed that we accurately predicted both uptake and outgassing,” says Glascoe. “We successfully developed an advanced predictive capability by executing a coordinated, simultaneous effort of all three aspects of our approach—experiments to characterize materials, model development, and validation experiments.”

From Weapons to Art

Accurate predictions of material compatibilities as a function of age are important in various fields, from designing aerospace components and medical devices to preserving works of art. Some chemical reactions between materials in sealed environments may be benign, but many of them will cause damage and loss of material functionality. Glascoe says, “Through this project, we have created an integrated, versatile capability to better analyze headspace gases in different weapons systems.” The team is now working with other Laboratory programs to improve the model and begin using the capability for real applications.

According to Glascoe, the success of this project is another example of what can be accomplished through a multidisciplinary approach to building capabilities. “At the Laboratory, we can collaborate with experts in many fields from physics to engineering to computational modeling,” she says. “Livermore also has the resources and a team approach that is critical to making key advances in many areas of national importance, including stockpile stewardship.”

—Caryn Meissner

Key Words: chemical kinetics, diffusion, nuclear weapon, PSUADE (Problem Solving Environment for Uncertainty Analysis and Design Exploration), reactive transport, sorption, stockpile stewardship, weapons aging.

For further information contact Elizabeth Glascoe (925) 424-5194 (glascoe2@llnl.gov).



An Increased Shot Rate at the National Ignition Facility

NESTLED against the eastern slopes of the Livermore Valley lies a one-of-a-kind engineering marvel that brings the cosmos to Earth. The National Ignition Facility (NIF) at Lawrence Livermore is a cornerstone of research that maintains the nation's nuclear weapons stockpile without full-scale testing. With its 192 laser beams, NIF can replicate the extreme conditions inside nuclear weapons, allowing scientists to examine thermonuclear burn in a controlled setting.

In 2012, an appropriations bill passed by Congress mandated that Livermore managers develop strategies to increase the number of NIF experiments, or shots, performed each year. In response, the Laboratory is updating the facility's operations and procedures to streamline the turnaround time between experiments. The increased shot rate will allow stockpile stewardship research to advance more quickly and will enhance NIF's utility as a uniquely

capable user facility for other national-security applications and high-energy-density (HED) science experiments.

The 120-Day Study's Recommendations

To understand how changes to NIF operations would affect research teams, the Laboratory conducted a series of meetings, conference calls, and site visits to discuss various options. These discussions, part of a 120-day study requested by the National Nuclear Security Administration, brought together 80 experts from Sandia and Los Alamos national laboratories, the University of Rochester, and other organizations involved in HED research. "The engagement from contributors was fabulous," says NIF Director Mark Herrmann. "I think it's a testament to how much the community cares about NIF's evolution. People really took the time to provide valuable feedback."

Systems operator Erik Mertens installs a fuel reservoir prior to a cryogenic layered ignition shot at the National Ignition Facility (NIF). A design change using an external fuel reservoir reduces the effort required to prepare the target positioner for these shots.

The study highlighted several opportunities to increase efficiency throughout the lifecycle of a NIF experiment. One major change targeted the facility's scheduling approach. "We used to schedule experiments on a shot-by-shot basis, which meant that project teams tried to gather as much information as they could from a single shot," says Doug Larson, NIF facility manager and a coleader for the 120-day study. "Users wanted to maximize the amount of data they retrieved in every experiment, which meant that designing a single shot and configuring its setup were time- and work-intensive." Researchers at the University of Rochester's Laboratory for Laser Energetics recommended that instead of specifying a number of shots for each project, schedulers allot time to a research team. Such an approach motivates team members to design experiments to maximize the scientific output for a given block of time.

Peg Folta, acting director of the NIF User Office, points to this new scheduling philosophy as a major contributor to the increased shot rate. "The new approach gives teams incentives to determine which shots they should match together to optimize their allotted time." Since different shot configurations demand changes to NIF's calibrations, pairing shots with similar configurations improves efficiency, which equates to more experiments within the set time.

Improving Relations with Users


Managers at NIF are also streamlining review procedures for experiments. Because of the facility's specialized equipment, experiments must be designed to yield the desired data without harming NIF components and diagnostics. "We identified areas where we could make it easier for scientists to work with the facility without reducing the quality of the experiments," says Larson. Expert groups were established to provide researchers with iterative feedback about target materials, laser energy, pulse shapes, and diagnostic instruments. In addition, the NIF team assembled recommended platforms that researchers can use to scaffold experiments. By implementing one of the standard platforms, project teams sidestep the time-consuming process of evaluating configurations for their experiments.

Workflow tools were also developed to guide teams throughout an experiment, from proposal submission and project reviews to experiment scheduling and data analysis. "We evaluated the entire experimental lifecycle to improve day-to-day activities for everyone," notes Folta.

Online user forums were launched so researchers can ask questions and provide feedback to each other. NIF managers also communicate administrative decisions via the forums. "We're increasing the visibility of our choices about NIF's new capabilities," says Larson. "We want to improve transparency so users understand why we prioritize certain capabilities at certain times." To date, NIF has added enhanced diagnostic tools, equipment, target materials, and laser alignments. "We're also building a third target positioner," says Larson, noting that another retractable arm for holding targets in place will allow for even faster shot turnaround.

Increased Efficiency

In addition to having a third target positioner, the 120-day study team identified a need for more efficient target fabrication. "More shots means more targets," explains Larson. Experimental success depends on the details of a target's design and fabrication. As researchers learn from their experiments, they will want not only more targets but also targets with a greater variety of features. The target fabrication team, which includes colleagues from other



At NIF, shot rates continue to increase. The facility is on target to achieve its goal of 300 shots in fiscal year (FY) 2015. Colors indicate the area being studied: high-energy-density (HED) physics, inertial confinement fusion (ICF), national security applications (NSA), discovery science (DS), and facility-related testing.

national laboratories and industrial partners General Atomics and Schafer Corporation, has implemented improvements to increase the efficiency of target design and production.

Designs vary from highly complex and precise cryogenic targets to those for experiments at room temperature. The most complex cryogenic targets require a layer of hydrogen “ice,” which must be grown on a target prior to an experiment. This process takes several days and ties up one of NIF’s two target positioners. Once the third positioner becomes operational, researchers will be able to perform successive room-temperature and cryogenic tests on two alternating arms while waiting for ice to form on targets held by the third positioner.

New facility maintenance procedures also improve efficiency. In conjunction with time allotments for experiments, NIF’s managers schedule two-day maintenance windows each week with specialized teams assigned to maintenance operations. This decision to designate the facility’s focus as “shot mode” or “maintenance mode” mirrors the approach used by the Laboratory for Laser Energetics and by the U.S. military.

A markedly successful innovation was to adopt the Formula One model for the NIF operations team. The team now follows task-driven staging and implementation procedures in much the same way that pit crews work at Formula One car races. With this approach, the NIF team can prepare for the next scheduled



Transport technicians remove the final optics damage inspection system from NIF’s target chamber. This work originally required nine technicians and took six hours. With an engineered lifting system, three technicians can remove the system in less than two hours. (Photograph by James Pryatel.)

shot while experiments are in progress. Operators can thus reconfigure the shot setup between experiments more quickly, which reduces laser downtime. Says NIF Operations Manager Bruno Van Wonterghem, “We’re beginning to see the effects of these improvements in early completion of the shot week. We’ve even had time to complete a few ‘opportunity shots’ beyond those originally planned for the week.”

On Target for Success

According to Larson, NIF is on schedule to complete 300 shots in fiscal year (FY) 2015. With the capabilities to be added in the coming year, NIF will likely achieve 400 shots in FY 2016—roughly double the number in FY 2013. Larson notes that the entire staff at NIF is proud of these accomplishments and continues to search for new efficiencies. “We’re pulling out all the stops to increase the value of the facility for our sponsors and to better support stockpile stewardship,” he says.

Mark Herrmann agrees. “NIF can create the extreme pressures and temperatures required to examine the operation of modern nuclear weapons,” he says. “By doing more experiments, researchers can learn at a faster rate and apply those lessons to maintaining the nation’s nuclear deterrent. NIF users working on other national-security applications and HED science experiments will also benefit from high shot rates.”

—M. H. Rubin

Key Words: 120-day study, high-energy-density (HED) physics, National Ignition Facility (NIF), nuclear weapons, operational efficiency, stockpile stewardship.

For further information contact Doug Larson (925) 422-1524 (larson8@llnl.gov).



A target area operator wearing personal protective equipment checks the installation of a diagnostic cart for an upcoming NIF experiment. Reconfiguring the target diagnostics is a critical step in setting up experiments. Streamlined operations have helped reduce the turnaround time between shots. (Photograph by James Pryatel.)

In this section, we list recent patents issued to and awards received by Laboratory employees. Our goal is to showcase the distinguished scientific and technical achievements of our employees as well as to indicate the scale and scope of the work done at the Laboratory. For the full text of a patent, enter the seven-digit number in the search box at the U.S. Patent and Trademark Office's website (<http://www.uspto.gov>).

Patents

Method and System for Compact, Multi-Pass Pulsed Laser Amplifier

Alvin Charles Erlandson

U.S. Patent 8,896,913 B1

November 25, 2014

Systems and Methods for Generation of Hydrogen Peroxide Vapor

Adam H. Love, Joel Del Eckels, Alexander K. Vu, Armando Alcaraz, John G. Reynolds

U.S. Patent 8,899,556 B2

December 2, 2014

Three Dimensional Microelectrode System for Dielectrophoresis

Dietrich A. Dehlinger, Klint A. Rose, Maxim Shusteff, Christopher G. Bailey, Raymond P. Mariella, Jr.

U.S. Patent 8,900,434 B2

December 2, 2014

Room-Temperature Quantum Noise Limited Spectrometry and Methods of the Same

Charles G. Stevens, Joseph W. Tringe

U.S. Patent 8,901,495 B2

December 2, 2014

Spatial Filters for High Power Lasers

Alvin Charles Erlandson, Andrew James Bayramian

U.S. Patent 8,902,497 B2

December 2, 2014

Nanolipoprotein Particles and Related Methods and Systems for Protein Capture, Solubilization, and/or Purification

Brett A. Chromy, Paul Henderson, Paul D. Hoeprich, Jr.

U.S. Patent 8,907,061 B2

December 9, 2014

Position Sensor for Linear Synchronous Motors Employing Halbach Arrays

Richard Freeman Post

U.S. Patent 8,917,086 B2

December 23, 2014

Awards

The **American Physical Society (APS)** has selected ten Laboratory scientists as **APS Fellows** in 2014.

Physicist **Michael Armstrong** was cited by the Topical Group on Instrument and Measurement Science for “outstanding contributions to time-domain experimental methods applied to materials under extreme conditions.”

Chris Barty, chief technology officer for the National Ignition Facility and Photon Science Principal Directorate, was recognized by the Division of Laser Science for “pioneering contributions to the advancement of ultrahigh-intensity laser science and to the development of laser-based X-ray and gamma-ray science.”

Physicist **Ray Beach** was acclaimed by the Division of Laser Science for “seminal contributions to high-average-power diode-end-pumped lasers, including many breakthroughs, widely adopted by the laser community, that have helped push such lasers to higher average powers and efficiencies, and for leadership in developing diode-pumped alkali-vapor lasers, and models for coherent and incoherent photon echoes.”

Physicist **Debbie Callahan** was honored by the Division of Plasma Physics for “innovative design and modeling of hohlraums for inertial confinement fusion and leadership in the execution of hohlraum experiments on the National Ignition Facility.”

Tony Gonis, an expert in theoretical solid-state physics, was acknowledged by the Division of Computational Physics for “advancing multiple scattering theory electronic structure methods

for metals, alloys and interfaces and for the dissemination of these techniques in condensed matter and materials science.”

Physicist **Frederic Hartemann** was selected by the Division of Physics of Beams for “remarkable insights and significant contributions to the physics of coherent radiation interacting with relativistic electrons.”

Physicist **Nobuhiko Izumi** was cited by the Topical Group on Instrument and Measurement Science for “outstanding contributions to the development of novel neutron and X-ray diagnostic capabilities for inertial confinement fusion experiments.”

Robert Rudd, group leader for Computational Materials Science in the Condensed Matter and Materials Division, was cited by the Division of Computational Physics for “seminal contributions to multiscale modeling of materials physics and science in support of national security.”

Scientist **James Tobin** was recognized by the Division of Condensed Matter Physics for the “use of soft X-ray spectroscopy to investigate complex systems, including actinide-based materials.”

Scientist **Yinmin (Morris) Wang** was honored by the Division of Materials Physics for “his major contributions to the understanding of deformation physics of nanocrystalline and nanotwinned materials, and for developing effective strategies to enhance the ductility of these superstrong materials for technological applications, including fusion energy targets.”

Awards

Each year, no more than one-half of 1 percent of the current APS membership is recognized by their peers through election to the status of fellow. APS fellowship recognizes members who have made advances in knowledge through original research and publication or those who have made significant and innovative contributions in the application of physics to science and technology. APS fellows also may have made significant contributions to the teaching of physics or service and participation in the activities of the society.

Lawrence Livermore National Laboratory has received a **Sustainability Award** from the **National Nuclear Security Administration (NNSA)** in the water category of the award's Environmental Stewardship division. NNSA gives 15 Sustainability Awards each year to its national labs and sites to recognize exemplary individual and team performance in advancing sustainability objectives through innovative and effective programs and projects that increase energy, water, and fleet efficiency and reduce greenhouse gases, pollution, and waste. The Laboratory was recognized for its water conservation plan.

Former Laboratory engineering associate director **Steve Patterson** has received the **2014 Lifetime Achievement Award** from the **American Society for Precision Engineering**. The lifetime achievement award is designated to those individuals who, over the span of their careers, made a significant impact to the science and discipline of precision engineering. Patterson received the award for his contributions to precision machine design, diamond turning machine control, and ultraprecision dimensional measurement.

Four former Lawrence Livermore researchers—**Martin Casado**, **Bill Colston**, **Fred Milanovich**, and **David Tuckerman**—were inducted into the Lawrence Livermore's **Entrepreneurs' Hall of Fame (EHF)**. The researchers, who represent the second "class" of inductees into the Laboratory's EHF, were honored for developing technologies during or after their Laboratory careers that created major economic impacts and spawned influential companies.

Collectively, this year's four inductees started three different companies. Casado and his team founded Nicira Networks, a company that developed computer applications for software-defined networking and network virtualization. In 2012, the firm was purchased by VMware.

In 2008, Bill Colston and Fred Milanovich founded QuantaLife, Inc., a company focused on deploying sensitive and accurate genetic testing technology. In 2011, Bio-Rad Laboratories Inc., a biomedical technology company, purchased QuantaLife.

Tuckerman cofounded nCHIP, Inc., in 1989, a technology company that developed enhanced microchips and multichip module systems. In 1995, nCHIP was purchased by Flextronics International Ltd.

"These kinds of contributions and success stories are one of the best measures of the innovation that comes out of the Laboratory and the way we approach our mission with an eye to innovation," says Laboratory Director Bill Goldstein. Livermore established the EHF to recognize current or former employees who have made major contributions to the United States through their inventiveness and entrepreneurial work in and with the private sector.

Matthew Levy, a Lawrence scholar in the Physics Division at Lawrence Livermore, has been awarded the prestigious **Newton International Fellowship** by the **Royal Society of the United Kingdom**. "Matthew has been identified as one of America's brightest theoretical plasma physicists to have graduated in the past decade," says Peter Norreys, Levy's sponsor and professor of inertial fusion science at the University of Oxford and Plasma Physics group leader at the Central Laser Facility, Rutherford Appleton Laboratory.


The highly competitive fellowship program makes 40 awards per year across all disciplines of the sciences and humanities, providing the opportunity for the best early-stage postdoctoral researchers to work at United Kingdom research institutions for a two-year period. Levy is the first American physicist to become a Newton fellow and will carry out his research at the University of Oxford.

A Hub for Collaborative Innovation

Modeled after research parks, the Livermore Valley Open Campus (LVOC) is an innovation hub for unclassified research and development conducted in partnership with private companies, academia, and a wide range of other organizations. Livermore's "anchor tenant" at LVOC is the High Performance Computing Innovation Center (HPCIC). Here, access to some of the world's most powerful supercomputers is accelerating the technological innovation that underpins America's economic vitality. In bioscience, energy resources, "big data" analysis, and other fields, HPCIC enables both short-term projects and long-term strategic partnerships that benefit the U.S. economy and further stockpile stewardship and other Laboratory missions.

Contact: Frederick Streitz (925) 423-3236 (streitz1@llnl.gov).

Helping Cities Prepare for a Disaster





Advanced modeling shows how urban dwellers can survive both natural and human-caused events.

Also in April/May

- Aerogel "glue" enables researchers to create three-dimensional carbon structures with superior properties for energy applications.
- A new electrophoretic deposition technique for additive manufacturing uses photoconductive electrodes to build three-dimensional multimaterial composites.
- Simulations of how water, sunlight, and semiconductor materials interact provide insights on the formation of hydrogen fuel.

Science & Technology Review
Lawrence Livermore National Laboratory
P.O. Box 808, L-664
Livermore, California 94551

PRSR STD
U.S. POSTAGE
PAID
San Bernardino, CA
PERMIT NO. 3330



Printed on recycled paper.